arXiv:1507.06240v3 [cs.DS] 20 Jul 2016

Sublinear-Space Distance Labeling using Hubs

Paweł Gawrychowski[1], Adrian Kosowski[2], and Przemysław Uznański[3]

[1]Institute of Informatics, University of Warsaw, Poland
[2]Inria Paris and IRIF, Université Paris Diderot, France
[3]Department of Computer Science, ETH Zürich, Switzerland

Abstract

A distance labeling scheme is an assignment of bit-labels to the vertices of an undirected, unweighted graph such that the distance between any pair of vertices can be decoded solely from their labels. We propose a series of new labeling schemes within the framework of so-called hub labeling (HL, also known as landmark labeling or 2-hop-cover labeling), in which each node u stores its distance to all nodes from an appropriately chosen set of hubs $S(u) \subseteq V$. For a queried pair of nodes (u, v), the length of a shortest $u-v$-path passing through a hub node from $S(u) \cap S(v)$ is then used as an upper bound on the distance between u and v.

We present a hub labeling which allows us to decode exact distances in sparse graphs using labels of size sublinear in the number of nodes. For graphs with at most n nodes and average degree Δ, the tradeoff between label bit size L and query decoding time T for our approach is given by $L = \mathcal{O}(n \log\log_\Delta T / \log_\Delta T)$, for any $T \leq n$. Our simple approach is thus the first sublinear-space distance labeling for sparse graphs that simultaneously admits small decoding time (for constant Δ, we can achieve any $T = \omega(1)$ while maintaining $L = o(n)$), and it also provides an improvement in terms of label size with respect to previous slower approaches.

By using similar techniques, we then present a 2-additive labeling scheme for general graphs, i.e., one in which the decoder provides a 2-additive-approximation of the distance between any pair of nodes. We achieve almost the same label size-time tradeoff $L = \mathcal{O}(n \log^2 \log T / \log T)$, for any $T \leq n$. To our knowledge, this is the first additive scheme with constant absolute error to use labels of sublinear size. The corresponding decoding time is then small (any $T = \omega(1)$ is sufficient).

We believe all of our techniques are of independent value and provide a desirable simplification of previous approaches.

1 Introduction

Distance labeling schemes, popularized by Gavoille et al. [16], are among the most fundamental distributed data structures for graph data. The design problem combines two major challenges. First of all, distance labelings serve the role of a *distance oracle*, i.e., a data structure which for a given undirected graph $G = (V, E)$ can answer queries of the form: "what is the distance between the nodes $s, t \in V$?". Throughout most of this paper, we will assume that G is an unweighted graph with n nodes and m edges. The efficiency of a distance oracle is measured by the interplay between the *space* requirement of the data structure representation, the *encoding time* required to set up the oracle for a given graph, and perhaps more importantly, its *decoding time*, that is, the time of processing a $s - t$ distance query. Moreover, a distance labeling scheme is defined more restrictively than a distance oracle, as an assignment of a binary string (label) $\mathsf{label}(u)$ to each node $u \in V$, so that the graph distance between u and v is uniquely determined by the pair of labels: $\mathsf{label}(u)$ and $\mathsf{label}(v)$. The *size* of a distance labeling scheme is now the maximum length of a node label in the graph. In this way, distance labelings add an extra layer of complexity to the graph distance decoding problem, by imposing a distributed representation of information in the labels $(\mathsf{label}(u) : u \in V)$. Whereas the concatenation of all n labels in a distance labeling forms a centralized distance oracle, distance labelings can also be applied in a distributed setting, in which the label of each node is stored at a distinct location in the network. This is the case, for instance, in applications in compact routing protocols, where the goal is to find a shortest path from a source node to a target node with a known label [12].

An interesting characteristic of the problem of distance oracle design for sparse graph is its inherent link to an underlying set intersection task. On the side of lower bounds, this is most clearly observed, following Pătraşcu and Roditty [21], when we consider a pair of vertices belonging to the same partition of a bipartite graph. The distance between them is 2 if and only if the sets of their neighbors intersect, and at least 4 otherwise. Consequently, assuming a plausible conjecture on the space required to decide intersection of a set of small sets, it follows that any oracle for graphs with $\widetilde{\mathcal{O}}(1)$ maximum degree, which admits constant decoding time, requires $\widetilde{\Omega}(n^2)$ space. (Here, the $\widetilde{O}$ and $\widetilde{\Omega}$ notation disregards polylogarithmic factors in n.) By contrast, many efficient algorithms for answering distance queries in real-world scenarios rely on the premise that the distance between a pair of nodes can be computed using an intersection-type query on a pair of small sets. In the basic framework of hub labelings, see [1], (introduced in [13] under the name of 2-hop covers, and also referred to as landmark labelings [3]), each node $u \in U$ stores the set of its distances to some subset $S(u) \subseteq V$ of other nodes of the graph. Then, the computed distance value $\delta'(u, v)$ for a queried pair of nodes $u, v \in V$ is returned as:

$$\delta'(u, v) := \min_{w \in S(u) \cap S(v)} \delta(u, w) + \delta(w, v), \tag{1}$$

where δ denotes the shortest path distance function between a pair of nodes. The computed distance between all pairs of nodes u and v is exact if set $S(u) \cap S(v)$ contains at least one node on some shortest $u - v$ path. This property of the family of sets $(S(u) : u \in V)$ is known as *shortest path cover*. The hub-based method of distance computation is in practice effective for two reasons. First of all, for transportation-type networks it is possible to show bounds on the sizes of sets S, which follow from the network structure. Notably, Abraham et al. [2] introduce the notion of highway dimension h of a network, which is presumed to be a small constant e.g. for road networks, and show that an appropriate cover of all shortest paths in the graph can be achieved using sets S of size $\widetilde{\mathcal{O}}(h)$. Moreover, the order in which elements of sets $S(u)$ and $S(v)$ is browsed when performing the minimum operation is relevant, and in some schemes, the operation can be interrupted once it is certain that the minimum has been found, before probing all elements of the set. This is the principle of numerous heuristics for the exact shortest-path problem, such as contraction hierarchies and algorithms with arc flags [9, 18].

In this work, we make use of the hub set techniques to obtain better (distributed) distance labelings. Whereas $\Omega(n)$ is a lower bound of the size of a hub set for general graphs, we provide hub-based schemes using smaller sets for specific case, leading to labels which can be encoded on $o(n)$ bits. Our scheme provides a shortest-path cover in the class of sparse graphs (with average degree $\Delta = 2m/n$ subpolynomial in n). This construction is overviewed in more detail in Section 1.2.

The implications of our result can be seen as twofold. First of all, our approach directly leads to labeling of smaller size (and smaller decoding time) for exact distance queries in sparse graphs than all previous distance labeling approaches. Additionally, an corollary of our result concerns the case of *k-additive* approximate distance labeling, in which the distance decoder is required to return an upper bound on the shortest path length which is within an additive factor of at most k from the optimum. So far, no way to construct k-additive distance labels using labels of sublinear size in n was known for any constant $k > 0$. (This question was considered previously in, e.g., [7]). We provide a way to construct a 2-additive distance labeling in general graphs using distance labels of size $o(n)$. (This result is essentially the best possible, since a 1-additive distance labeling requires distance labels of size at least $n/4$ already on the class of bipartite graphs.)

In our approaches, the size of the obtained distance labels for the considered cases is improved with respect to the state-of-the-art by up to a logarithmic multiplicative factor. Rather than seeing this as "gaining" a logarithm, we rather see this as "not losing" a logarithm. Indeed, the basic ingredient of the hub sets in previous approaches was a subset of nodes, sampled independently at random from V [6, 10]. The constructions then relied on the probabilistic method to guarantee that the hubs would have the shortest-path cover properties, based on the premise that for each pair of nodes the constructed hubs provide a shortest path cover with sufficiently high probability. The derandomization of this process resulted in a loss of a logarithmic factor in the analysis of the size of labels. Our approach shows how to avoid this issue: when constructing labelings for sparse graphs, we do away with randomization altogether, relying on simple structural results to replace the random subset of nodes.

1.1 Related Work

Distance Labelings. The distance labeling problem in undirected graphs was first investigated by Graham and Pollak [17], who provided the first labeling scheme with labels of size $\mathcal{O}(n)$. The decoding time for labels of size $\mathcal{O}(n)$ was subsequently improved to $\mathcal{O}(\log \log n)$ by Gavoille et al. [16] and to $\mathcal{O}(\log^* n)$ by Weimann and Peleg [25]. Finally, Alstrup et al. [7] present a scheme for general graphs with decoding in $\mathcal{O}(1)$ time using labels of size $\frac{\log 3}{2}n + o(n)$ bits.* This matches up to low order terms the space of the currently best known distance oracle with $\mathcal{O}(1)$ time and $\frac{\log 3}{2}n^2 + o(n^2)$ total space in a *centralized* memory model, due to Nitto and Venturini [19].

The notion of D-preserving distance labeling, first introduced by Bollobás et al. [10], describes a labeling scheme correctly encoding every distance that is at least D. [10] presents such a D-preserving scheme of size $\mathcal{O}(\frac{n}{D}\log^2 n)$. This was recently improved by Alstrup et al. [6] to a D-preserving scheme of size $\mathcal{O}(\frac{n}{D}\log^2 D)$. Together with an observation that all distances smaller than D can be stored directly, this results in a labeling scheme of size $\mathcal{O}(\frac{n}{R}\log^2 R)$, where $R = \frac{\log n}{\log \frac{m+n}{n}}$. For sparse graphs, this is $o(n)$.

For specific classes of graphs, Gavoille et al. [16] described a $\mathcal{O}(\sqrt{n}\log n)$ distance labeling for planar graphs, together with $\Omega(n^{1/3})$ lower bound for the same class of graphs. Additionally, $\mathcal{O}(\log^2 n)$ upper bound for trees and $\Omega(\sqrt{n})$ lower bound for sparse graphs were given.

Distance Labeling with Hub Sets. For a given graph G, the computational task of minimizing the sizes of hub sets $(S(u) : u \in V)$ for exact distance decoding is relatively well understood. A $\mathcal{O}(\log n)$-approximation algorithm for minimizing the average size of a hub set having the sought shortest path cover property was presented in Cohen et al. [13], whereas a $\mathcal{O}(\log n)$-approximation for minimizing the largest hub set at a node was given more recently in Babenko et al. [8]. Rather surprisingly, the structural question of obtaining bounds on the size of such hub sets for specific graph classes is wide open. For example, for the class of graphs of constant maximum degree, there is a large gap between the hub sets in our construction (of size $\mathcal{O}(n/\log n)$) and the generic lower bound of $\widetilde{\Omega}(\sqrt{n})$.

Distance Oracles. A centralized version of distance labeling problem is *distance oracle* problem, where one asks for a centralized data structure allowing for querying a distance between pair of vertices. There

*For the sake of sanity of the notation, we define $\log x = \max(1, \log_2(x))$.

usually one asks for what type of tradeoffs are possible between size of the structure, time of the query and allowed error (multiplicative stretch). Sommer et al. [23] proved that any constant time, constant stretch oracle must be superlinear in n. Thorup and Zwick [24] proved that distance oracles of stretch 2 require $\Omega(n^2)$ space, and of stretch 3 require $\Omega(n^{3/2})$ space. Pătraşcu and Roditty [21] strengthened the lower bound for stretch 2, proving a lower bound of $\Omega(n\sqrt{m})$ on the size of oracles with constant query time. For general weighted graphs, Thorup and Zwick [24] designed a distance oracle of size $\mathcal{O}(kn^{1+1/k})$, stretch-$(2k-1)$ and $\mathcal{O}(k)$ time. The query time has been improved to $\mathcal{O}(\log k)$ time by Wulff-Nilsen [26], and to constant time in Chechik [11]. The size of the distance oracle from [24] is optimal assuming girth-conjecture. For sparse graphs, [21] design distance oracle of size $\mathcal{O}(n^{4/3}m^{1/3})$ and stretch 2 in constant time. Also in [21], a conditional lower bound of $\widetilde{\Omega}(n^2)$ bits for a constant time distance oracle is provided. Cohen and Porat [14] extended this result to sparse graphs. An up-to-date survey of results on approximate distance oracles is provided in [22].

1.2 Our Results and Organization of the Paper

We start by introducing the necessary conventions in Section 2. We also describe the basic building block for encoding distance labels, namely, an efficient method of storing the hub set of a node, together with corresponding distances, in its distance label.

In Section 3, we show how to construct an exact distance labeling scheme for graphs of bounded maximum degree. This relies on hub sets which consist, for a vertex u of the union of all nodes from a small ball around vertex u, and all nodes from a selection of equally-spaced levels of the breadth-first-search tree of u. We then apply a trick, known from the previous work of [4], to reduce the problem of constructing a labeling scheme for a graph with bounded average degree to that of constructing a labeling scheme for a bounded-degree graph on twice as many nodes. For graphs with at most n nodes and average degree Δ, the tradeoff between label bit size L and query decoding time T for our approach is given by $L = \mathcal{O}(n \log\log_\Delta T / \log_\Delta T)$, for any $T \leq n$. In particular, setting $T = n$, we obtain labels of size $\mathcal{O}(\frac{n}{R} \log R)$, which improves previously best result [6] by a factor of $\log R$, keeping the $\widetilde{\mathcal{O}}(n)$ decoding time. On the other end, setting $T = \log n$ we obtain first sublinear size distance labeling that achieves almost-constant decoding time.

In Section 4, we adapt our approach to general graphs, using a variant of the proposed labeling scheme for sparse graphs to achieve 2-additive approximation of distances. As before, we achieve a tradeoff between label size L and time T of the form $L = \mathcal{O}(n \log^2 \log T / \log T)$, for any $T \leq n$. This 2-additive distance labeling scheme can be easily transformed into an exact one, by encoding the difference between the estimation and the true distances. Since this difference is always from $\{0, 1, 2\}$, we achieve labels of size $\frac{\log 3}{2} n + o(n)$ (with any $\omega(1)$ decoding time), or of size $(\frac{\log 3}{2} + \varepsilon)n$ (with $\mathcal{O}(1)$ decoding time), for any $\varepsilon > 0$. Our approach almost matches the size of the best known distance labeling schemes [7], which make use of labels of size $\frac{\log 3}{2} n + o(n)$ to achieve $\mathcal{O}(1)$ decoding time. Arguably, our approach may be considered simpler.

We remark that all our results apply to unweighted graphs, in which each edge has unit length. For sparse graphs, in which each edge has an integer weight from some interval $[1, W]$, we can use the same hub sets with an appropriately modified encoding to achieve a time-label tradeoff of $L = \mathcal{O}(n \log\log_\Delta T \log W / \log_\Delta T)$. For the additive scheme, by subdividing each edge of length $w \in [1, W]$ into a chain of unweighted edges (of length 1), we achieve a conversion of the 2-additive distance labeling scheme into a $(2W)$-additive-distance scheme for weighted graphs.

2 Preliminaries

Notation and Conventions. Even though we are mainly interested in unweighted graphs, for technical reasons in Sections 3 and 4 we will work in a more general setting where every edge of a graph has a fixed cost from the set $\{0, 1\}$. $\delta(u, v)$ denotes the cost of a cheapest path connecting a pair of nodes u and v, and $\ell(u, v)$ denotes the smallest number of edges on such a path. We will require the constructed distance labeling to return the value of $\delta(u, v)$. The degree of a node v is denoted by $\deg(v)$. When analyzing the complexity

of the decoding, we assume standard word RAM with logarithmic word size, where we are allowed to access $\log n$ consecutive bits of the stored binary string in constant time.

From now on, we assume that the graph is connected. This is enough because we can always include the identifier of its connected component in the label of every node, and return ∞ if u and v belong to different connected components; this only induces additive $\log n$ overhead to the label size.

Encoding Distances and Identifiers. The basic procedure for encoding a hub set in a label exploits some ideas from [7]; we provide a self-contained exposition for completeness. We fix an arbitrary spanning tree of the graph and assign preorder numbers in the tree to the nodes, i.e., node numbered 1 corresponds to the root and so on. The preorder number of a node u is denoted by $\mathsf{name}(u)$. Such a numbering has the following useful property.

Lemma 2.1. *Let $v_1, v_2, \ldots, v_n$ be the preorder sequence of all nodes. Then, for any node u, $\sum_{i=2}^{n} |\delta(u, v_{i-1}) - \delta(u, v_i)| \leq 2n$.*

Proof. Consider an Euler tour corresponding a traversal of the chosen spanning tree. Every node is visited at least once there, and the total length of the tour is at most $2n$. Consequently, we can cut the tour into paths connecting node v_{i-1} with node v_i, for every $i = 2, 3, \ldots, n$. The total length of all these paths is at most $2n$ and the claim follows. □

The following lemma is used for encoding a hub set S using $\mathcal{O}(|S| \log(n/|S|))$ bits.

Lemma 2.2. *For a fixed v and set S such that $|S| \leq \frac{n}{x}$, set S and all of the distances $\delta(v, u)$ for $u \in S$ can be stored in $\mathcal{O}(\frac{n}{x} \log x)$ bits. For any constant $t > 0$, the representation can be augmented with $\mathcal{O}(\frac{n}{\log^t n})$ additional bits so that all elements of S can be extracted one-by-one in $\mathcal{O}(|S|)$ total time and given any u we can check if $u \in S$ (and if so, extract $\delta(u, v)$) in $\mathcal{O}(1)$ time.*

Proof. Let $S = (v_1, \ldots, v_{|S|})$, where $\mathsf{name}(v_1) < \mathsf{name}(v_2) < \ldots < \mathsf{name}(v_{|S|})$. We store $\mathsf{name}(v_1)$ and then the differences $\mathsf{name}(v_2) - \mathsf{name}(v_1), \ldots, \mathsf{name}(v_{|S|}) - \mathsf{name}(v_{|S|-1})$. Every difference is encoded using the Elias γ code (see Elias [15]), and the encodings are concatenated to form one binary string. We are storing up to $\frac{n}{x}$ integers whose absolute values sum up to at most n, so by Jensen's inequality this takes $\mathcal{O}(\frac{n}{x} \log x)$ bits in total. Similarly, we store $\delta(u, v_1)$ and then the differences $\delta(u, v_2) - \delta(u, v_1), \ldots, \delta(u, v_{|S|}) - \delta(u, v_{|S|-1})$. By Lemma 2.1 we are again storing up to $\frac{n}{x}$ numbers whose absolute values sum up to at most $2n$, which takes $\mathcal{O}(\frac{n}{x} \log x)$ bits.

All v_i can be extracted one-by-one in $\mathcal{O}(1)$ time each with standard bitwise operations. To facilitate checking if $x \in S$ in $\mathcal{O}(1)$ time, we observe that it is enough to store a bit-vector $B[1..n]$, where the $\mathsf{name}(v_i)$-th bit is set to $\mathbf{1}$, for every $i = 1, 2, \ldots, |S|$. Then checking if $x \in S$ reduces to two $\mathsf{rank}_\mathbf{1}$ queries. A $\mathsf{rank}_\mathbf{1}$ query counts $\mathbf{1}$s in the specified prefix of the bit-vector and a $\mathsf{select}_\mathbf{1}$ query returns the position of the k-th $\mathbf{1}$ in the bit-vector. By the result of Pătraşcu [20], for any constant $t > 0$, a bit-vector of length n containing $\frac{n}{x}$ $\mathbf{1}$s can be stored using

$$\log \binom{n}{\frac{n}{x}} + \mathcal{O}(\frac{n}{\log^t n}) = \mathcal{O}(\frac{n}{x} \log x) + \mathcal{O}(\frac{n}{\log^t n})$$

bits so that any rank or select query can be answered in $\mathcal{O}(t)$ time. This allows us to check if $u \in S$ and calculate i such that $u = v_i$ in $\mathcal{O}(t)$ time. To retrieve $\delta(u, v_i)$, we store two additional bit-vectors B_+ and B_-. Each of them contains exactly $\frac{n}{x}$ $\mathbf{1}$s and up to $2n$ $\mathbf{0}$s. The bit-vectors are defined as follows. For each $i = 2, 3, \ldots, n$ we consider the difference $\eta = \delta(u, v_i) - \delta(u, v_{i-1})$. If $\eta \geq 0$, we append $\mathbf{0}^\eta\mathbf{1}$ to B_+ and $\mathbf{1}$ to B_-. Otherwise, we append $\mathbf{1}$ to B_+ and $\mathbf{0}^{-\eta}\mathbf{1}$ to B_-. By Lemma 2.1, each of these two bit-vectors contains at most $2n$ $\mathbf{0}$s, so they can be stored using $\mathcal{O}(\frac{n}{x} \log x + \frac{n}{\log^t n})$ bits so that any rank or select query can be answered in $\mathcal{O}(t)$ time. To recover $\delta(u, v_i)$, we need to sum up all the differences. This reduces to summing up all positive and all negative differences separately, which can be done using the corresponding bit-vector with one $\mathsf{rank}_\mathbf{1}$ and one $\mathsf{select}_\mathbf{0}$ query in $\mathcal{O}(t)$ total time. □

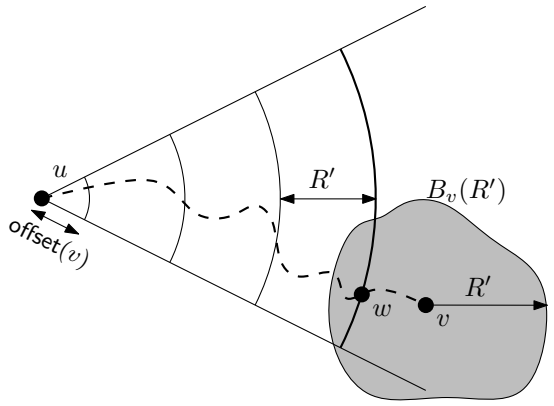

Figure 1: Shortest path from u to v goes through w which belongs to both $B_v(R')$ and $L_u(\mathsf{offset}(u))$.

We remark that the above encoding and decoding scheme is efficient for sets of size $|S| = \widetilde{\mathcal{O}}(n)$. For smaller sets, we will simply use an explicit encoding of all distances in S, requiring $\mathcal{O}(|S| \log n)$ bits.

3 Exact Distance Labeling in Sparse Graphs

3.1 Graphs of Bounded Maximum Degree

In this subsection, we assume that $\deg(u) \leq \Delta$ for every node u. We consider distance labeling schemes characterized by a time parameter T. Intuitively, in the construction, $R = \frac{\log T}{\log \Delta}$ will be a threshold parameter, distinguishing small distances from large distances in the graph — a node will be able to afford to explicitly store the distances and identifiers of all nodes up to some distance $\mathcal{O}(R)$ from itself in its distance label. Although this case is of independent interest, we are considering it as a building block for construction of labeling in graphs of bounded average degree. Thus graphs considered here are weighted with edge weights from $\{0, 1\}$, for the reason explained in Section 3.2.

The rest of this subsection is devoted to the proof of the following Theorem.

Theorem 3.1. *Fix any value $\Delta \leq T \leq n$ and let $R = \log_\Delta T$. In bounded-degree graphs, there is a labeling scheme of size $\mathcal{O}(\frac{n}{R} \log R)$ and decoding time $\mathcal{O}(T)$.*

Let us denote $R' = \lfloor R \rfloor$. Since $R' \geq 1$, we can bound $R \geq R' \geq \frac{1}{2}R$. Consider a node u. The ball of radius r centered at u, denoted $B_u(r)$, is the set of nodes which can be reached from u by following at most r edges. Because the degrees of all nodes are bounded by Δ, $|B_u(r)| = \mathcal{O}(\Delta^r)$. The k-th layer centered at u, denoted $L_u(k)$, consists of all nodes v such that $\ell(u, v) = k \pmod{R}'$. Because the layers are disjoint, there exists an $\mathsf{offset}(u) \in \{0, 1, \ldots, R' - 1\}$ such that $|L_u(\mathsf{offset}(u))| \leq \frac{n}{R'}$.

Definition of the Labeling. We define the hub set of node u, to which it stores all its distances, as $S(u) := B_u(R') \cup L_u(\mathsf{offset}(u))$, see Fig. 1. Formally, the label of u consists of the following:

1. n and $\mathsf{name}(u)$,
2. $\mathsf{name}(v)$ and $\delta(u, v)$ for every $v \in B_u(R')$,
3. $\mathsf{name}(v)$ and $\delta(u, v)$ for every $v \in L_u(\mathsf{offset}(u))$.

Computing $\delta(u,v)$. For reasons of efficiency, we will not perform the distance decoding following Eq. (1) directly, but we will treat the two components of the hub set of each node separately. Given $\mathsf{name}(u)$ and $\mathsf{name}(v)$, we can determine $\delta(u,v)$ as follows. First we check if $v \in B_u(R')$ and if so return the stored $\delta(u,v)$. Otherwise, we iterate through all nodes $w \in B_u(R')$ and check if $w \in L_v(\mathsf{offset}(v))$. If so, we know $\delta(u,w)+\delta(w,v)$. We return the smallest such sum.

For the proof of correctness of the distance decoder, it is clear that $\delta(u,w)+\delta(w,v) \geq \delta(u,v)$ for any w, so it remains to argue that either $v \in B_u(R')$ or there exists $w \in B_u(R')$ such that $w \in L_v(\mathsf{offset}(v))$ and $\delta(u,w)+\delta(w,v)=\delta(u,v)$. Consider a shortest path $(p_0,p_1,p_2,\dots p_\ell)$ where $v=p_0$ and $u=p_\ell$ such that $\ell=\ell(u,v)$. If $\ell \leq R'$, $v \in B_u(R')$ and there is nothing to prove, so we can assume that $\ell > R'$. Observe that for any $i=0,1,\dots,\ell$, $\ell(v,p_i)=i$, so in particular $p_{\alpha\cdot R'+\mathsf{offset}(v)} \in L_v(\mathsf{offset}(v))$ for any integer $\alpha \geq 0$. We choose $\alpha = \left\lfloor \frac{\ell-\mathsf{offset}(v)}{R'} \right\rfloor$ and $w = p_{\alpha\cdot R'+\mathsf{offset}(v)}$. Then $w \in L_v(\mathsf{offset}(v))$, $w \in B_u(R')$ by the choice of α, and $\delta(u,w)+\delta(w,v)=\delta(u,v)$ because w lies on a shortest path connecting u and v, so indeed we are able to correctly determine $\delta(u,v)$.

Encoding and Size of the Scheme. Encoding n and $\mathsf{name}(u)$ takes $\mathcal{O}(\log n)$ bits. The set $B_u(R')$ with corresponding distances is stored explicitly, while set $L_u(\mathsf{offset}(u))$ together with the corresponding distances is stored using Lemma 2.2, using $\mathcal{O}(\Delta^{R'}\log n)=\mathcal{O}(T\log n)$ and $\mathcal{O}(\frac{n}{R'}\log R')$ bits, respectively. Hence the total size of the scheme is

$$\mathcal{O}(\log n + T\log n + \frac{n}{R'}\log R') = \mathcal{O}(\frac{n}{R}\log R),$$

where we have used the fact that for any $T=\mathrm{poly}(n)$ the claimed label size is the same, thus we can assume $T=o(n/\mathrm{polylog}(n))$.

Complexity of the Decoding. Checking if $v \in B_u(R')$ and retrieving the encoded $\delta(u,v)$ takes $\mathcal{O}(T)$ time. Similarly, iterating through all $w \in B_u(R')$, checking if $w \in L_v(\mathsf{offset}(v))$ and if so retrieving the encoded $\delta(v,w)$ takes, by Lemma 2.2, $\mathcal{O}(1)$ time per single w, thus $\mathcal{O}(|B_u(R')|)=\mathcal{O}(T)$ total time. All in all, we can compute $\delta(u,v)$ in $\mathcal{O}(T)$ total time. □

Smaller values of T. For the sake of completeness, we consider the special case of $T<\Delta$. Consider labeling where the label of a node u consists of n, $\mathsf{name}(u)$, and all values $\delta(u,v)$ for $v \in V$ stored using Lemma 2.2. This takes $\mathcal{O}(n)$ bits, with $\mathcal{O}(1)$ decoding time, and matches claimed bounds from Theorem 3.1.

We also observe that our result applies not only to distance labels, but also as a size upper bound of hub sets for sparse graphs. Indeed, by fixing $T=n$, and observing that $|B_u(R')|+|L_u(\mathsf{offset}(u))| \leq \frac{n}{R'}$, we have the following:

Corollary 3.2. *In bounded-degree graphs, there is a hub set construction of size $\mathcal{O}(\frac{n}{\log_\Delta n})$ vertices per node.*

3.2 Graphs of Bounded Average Degree

We now allow for bounded average degree by reduction to the approach from Subsection 3.1. Given a graph G, let $\Delta=\frac{m+n}{n}$. We will create a new graph by splitting nodes of high degree. Following the formulation from [4, Lemma 4.2] (cf. Figure 2), we can obtain a graph G' on at most $2n$ nodes and at most $m+n$ edges, such that the degree of every node is bounded by $\lceil \frac{m}{n} \rceil + 2 \leq \Delta+2$ and the distance between two nodes in the original graph G is exactly the same as the distance between their corresponding nodes in the new graph G'. We can now directly apply the scheme from Theorem 3.1 to graph G', and exactly the same distance labels will work for the corresponding nodes of graph G. In this way, we obtain a scheme of size:

$$\mathcal{O}(\frac{n}{R}\log R + \frac{m}{n}) = \mathcal{O}(\frac{n}{R}\log R), \text{ where } R=\mathcal{O}(\frac{\log T}{\log \Delta}).$$

which returns $\delta(u,v)$ in $\mathcal{O}(T)$ time given the labels of u and v. The correctness of the this reduction is guaranteed by the fact that Theorem 3.1 allows for edge weights from $\{0,1\}$.

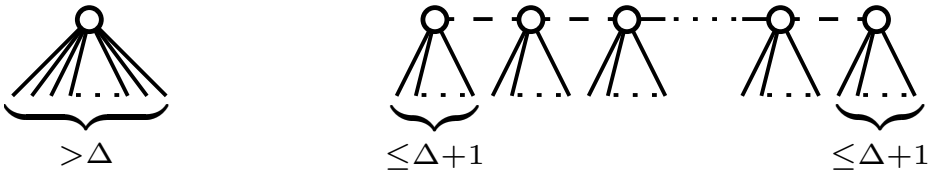

Figure 2: Example of subdividing of a large degree node (on the left) into a family of nodes of small degree, connected by edges of weight 0 (dashed edges).

Theorem 3.3. *Fix any $T \leq n$ and Δ, and let $R = \frac{\log T}{\log \Delta}$. There exists an exact distance labeling for graphs with average degree Δ using labels of size $\mathcal{O}(\frac{n}{R} \log R)$ and a corresponding decoding scheme requiring time $\mathcal{O}(T)$.* □

It is easy to see that this reduction preserves bounds on the size of hub sets, so we have the following:

Corollary 3.4. *In graphs with average degree Δ, there is a hub set construction of size $\mathcal{O}(\frac{n}{\log_\Delta n})$ vertices per node.*

4 2-Additive Distance Labeling in General Graphs

We will apply a similar distance labeling scheme as for sparse graphs, obtaining a 2-additive approximation of the distance between any pair of with label sizes of $o(n)$ per node. In this approximate scheme, the hub sets will have the following property. The label of each node $u \in V$ will provide an encoding of the node identifiers of a subset $S(u) \subseteq V$ and of the distances from u to all elements of $S(u)$. The sets $S(u)$ will be defined so that for any pair u, v, there exists a node $w \in S(u) \cap S(v)$, such that either w or a neighbor of w lies on the shortest path from u to v in G. We will decode the approximate distance as before, using Eq. (1); clearly, $\delta'(u, v) \in \delta(u, v) + \{0, 1, 2\}$.

The construction of sets $S(u)$ is performed as follows. Let $\tau < \frac{1}{2} \log n$ be an threshold value of vertex degree, to be chosen later. Let $V' = \{v \in V : \deg(v) > \tau\}$, and let $S' \subseteq V$ be a minimal dominating set for V', i.e., a subset of V with the property: $\forall_{w \in V'} B_w(1) \cap S' \neq \emptyset$. By a straightforward application of the probabilistic method (cf. [5, proof of Theorem 1.2.2]), we have that there is S' such that $|S'| \leq \frac{1+\ln(\tau+1)}{\tau+1} n < \frac{2 \ln \tau}{\tau} n$, and it can be easily constructed in polynomial time (a deterministic construction by a folklore greedy algorithm gives set of size $\mathcal{O}(\frac{\ln \tau}{\tau} n)$). For every $u \in V$, we define $B'_u(r)$ as the set of nodes of the ball of radius r around u in the subgraph $G[V \setminus V']$. Finally, we define $R = \frac{\tau}{\log \tau}$ and let L_u, R', and $\mathsf{offset}(u)$ be defined as in Section 3.1, and let S'_u be a minimal subset of S' such that for every $w \in V'$ adjacent the boundary of $B'_u(R')$, i.e. $B_w(1) \cap B'_u(R') \neq \emptyset$, we have $B_w(1) \cap S'_u \neq \emptyset$. Such S'_u can be easily constructed in polynomial time, and moreover, since there are at most $\tau^{R'+1}$ vertices adjacent to the boundary, we have $|S'_u| = 2^{\mathcal{O}(\tau)}$.

The approximate distance label of u now consists of the following elements:

1. n and $\mathsf{name}(u)$,
2. $\mathsf{name}(v)$ and $\delta(u, v)$ for every $v \in B'_u(R')$,
3. $\mathsf{name}(v)$ and $\delta(u, v)$ for every $v \in L_u(\mathsf{offset}(u))$.
4. $\mathsf{name}(v)$ and $\delta(u, v)$ for every $v \in S'_u$,
5. $\mathsf{name}(v)$ and $\delta(u, v)$ for every $v \in S' \setminus S'_u$.

The separation of S' into S'_u and $S' \setminus S'_u$ in the label is done to allow efficient decoding.

Computing $\delta(u,v)$. To show the correctness of this approximate labeling scheme, fix a pair of vertices $u, v \in V$. If there exists a vertex $w \in B'_u(R')$ lying on a fixed shortest path P between u and v such that $w = v$ or $w \in L_v(\mathsf{offset}(v))$, then the labeling scheme finds the shortest path distance between u and v as in Section 3.1. Otherwise, let y be the nearest vertex to u lying in $P \setminus B'_u(R')$; it follows from the construction that $y \in V'$. Then, there exists $w \in B_y(1)$ such that $w \in S'_u \subseteq S'$. In this case, the distance $\delta(u,w)+\delta(v,w)$ is a 2-additive approximation of $\delta(u,v)$.

Size of the Scheme. The size of the label of a node u in the scheme can be bounded as follows: $|B'_u(R')| \leq \tau^{R'} \leq 2^{\tau}$, $|L_u(\mathsf{offset}(u))| \leq \frac{n}{R}$, $S' < \frac{2\ln\tau}{\tau}n$. Overall, the total size is $|B'_u(R') \cup L_u(\mathsf{offset}(u)) \cup S'| = \mathcal{O}(\frac{\log\tau}{\tau}n)$, thus using Lemma 2.2 to store the sets and the corresponding distances we obtain labels of size $\mathcal{O}(n\frac{\log^2\tau}{\tau})$.

Complexity of the Decoding. To perform the distance decoding, for a given pair $u, v \in V$, it suffices to minimize $\delta(u,w)+\delta(v,w)$ over all w belonging to $B'_u(R') \cup S'_u$ which are also encoded in the label of v. Hence, distance decoding is possible in time $2^{\mathcal{O}(\tau)}$. Overall, setting $T := 2^{\mathcal{O}(\tau)}$, we obtain the following main result of the section.

Theorem 4.1. *There is a 2-additive distance labeling scheme for general graphs, which achieves decoding time T using labels of size $\mathcal{O}(n\frac{\log^2\log T}{\log T})$, for any $T \leq n$.* □

Finally, we remark on some implications of our result. By a standard argument, converting a 2-additive approximate distance labeling into an exact one requires an additional label of size $\frac{\log_2 3}{2}n$ bits per node (and an additional $\mathcal{O}(\frac{n}{\log n})$ overhead in the space, which is negligible), with each node u encoding the difference between the approximate and real distance value, $\delta'(u,v) - \delta(u,v)$, for all $v \in \{(u+1) \bmod n, \ldots, (u + \lfloor\frac{n}{2}\rfloor) \bmod n\}$. The time overhead of the corresponding decoding is $\mathcal{O}(1)$. In an analogous manner, converting a 2-additive approximate distance labeling into an 1-additive approximate one requires an additional label of size $\frac{1}{2}n$ bits per node. Thus we convert our scheme into an exact distance labeling scheme or 1-additive scheme achieving T decoding time using labels of size respectively $\frac{\log_2 3}{2}n+\mathcal{O}(n\frac{\log^2\log T}{\log T})$ or $\frac{1}{2}n+\mathcal{O}(n\frac{\log^2\log T}{\log T})$, for any $T \leq n$.

Thus, setting τ as an arbitrarily small increasing function of n, for any desired decoding time $T = \omega(1)$ we can make use of labels of size $o(n)$, $\frac{1}{2}n + o(n)$ and $\frac{\log_2 3}{2}n + o(n)$ respectively for 2-additive, 1-additive and exact distances. Moreover, using this scheme, $\mathcal{O}(1)$ decoding time can be achieved for labels of size εn, $(\frac{1}{2}+\varepsilon)\cdot n$ and $(\frac{\log_2 3}{2}+\varepsilon)\cdot n$, for any absolute constant $\varepsilon > 0$.

While a slightly stronger in terms of decoding time schemes were presented in Alstrup et al. [7] (achieving $\mathcal{O}(1)$ decoding time and labels of size $\frac{\log_2 3}{2}n + o(n)$ and $\frac{1}{2}n + o(n)$ for exact and 1-additive distances), we believe that presented here schemes are of independent value due to the simplification of the construction.

Acknowledgments

Most of the work was done while PU was affiliated to Aalto University, Finland. Research partially supported by the National Science Centre, Poland - grant number 2015/17/B/ST6/01897.

References

[1] I. Abraham, D. Delling, A. V. Goldberg, and R. F. Werneck. Hierarchical hub labelings for shortest paths. In *Proceedings of the 20th Annual European Conference on Algorithms*, ESA'12, pages 24–35, Berlin, Heidelberg, 2012. Springer-Verlag.

[2] I. Abraham, A. Fiat, A. V. Goldberg, and R. F. Werneck. *Highway Dimension, Shortest Paths, and Provably Efficient Algorithms*, chapter 64, pages 782–793.

[3] I. Abraham and C. Gavoille. On approximate distance labels and routing schemes with affine stretch. In *In International Symposium on Distributed Computing (DISC)*, pages 404–415, 2011.

[4] R. Agarwal and P. B. Godfrey. Distance oracles for stretch less than 2. In S. Khanna, editor, *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013, New Orleans, Louisiana, USA, January 6-8, 2013*, pages 526–538. SIAM, 2013.

[5] N. Alon and J. H. Spencer. *The Probabilistic Method*. J. Wiley & Sons, New York, 2nd edition, 2000.

[6] S. Alstrup, S. Dahlgaard, M. Bæk Tejs Knudsen, and E. Porat. Sublinear distance labeling for sparse graphs. *CoRR*, abs/1507.02618, 2015.

[7] S. Alstrup, C. Gavoille, E. B. Halvorsen, and H. Petersen. Simpler, faster and shorter labels for distances in graphs. In R. Krauthgamer, editor, *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 338–350. SIAM, 2016.

[8] M. A. Babenko, A. V. Goldberg, A. Gupta, and V. Nagarajan. Algorithms for hub label optimization. In F. V. Fomin, R. Freivalds, M. Z. Kwiatkowska, and D. Peleg, editors, *Automata, Languages, and Programming - 40th International Colloquium, ICALP 2013, Riga, Latvia, July 8-12, 2013, Proceedings, Part I*, volume 7965 of *Lecture Notes in Computer Science*, pages 69–80. Springer, 2013.

[9] R. Bauer and D. Delling. SHARC: Fast and robust unidirectional routing. *J. Exp. Algorithmics*, 14:4:2.4–4:2.29, Jan. 2010.

[10] B. Bollobás, D. Coppersmith, and M. Elkin. Sparse distance preservers and additive spanners. *SIAM Journal on Discrete Mathematics*, 19(4):1029–1055, 2005.

[11] S. Chechik. Approximate distance oracles with constant query time. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, STOC '14, pages 654–663, New York, NY, USA, 2014. ACM.

[12] V. Chepoi, F. F. Dragan, B. Estellon, M. Habib, Y. Vaxès, and Y. Xiang. Additive spanners and distance and routing labeling schemes for hyperbolic graphs. *Algorithmica*, 62(3-4):713–732, 2012.

[13] E. Cohen, E. Halperin, H. Kaplan, and U. Zwick. Reachability and distance queries via 2-hop labels. *SIAM J. Comput.*, 32(5):1338–1355, May 2003.

[14] H. Cohen and E. Porat. On the hardness of distance oracle for sparse graph. *CoRR*, abs/1006.1117, 2010.

[15] P. Elias. Universal codeword sets and representations of the integers. *IEEE Transactions on Information Theory*, 21(2):194–203, Mar 1975.

[16] C. Gavoille, D. Peleg, S. Pérennes, and R. Raz. Distance labeling in graphs. *J. Algorithms*, 53(1):85–112, Oct. 2004.

[17] R. Graham and H. Pollak. On embedding graphs in squashed cubes. In Y. Alavi, D. Lick, and A. White, editors, *Graph Theory and Applications*, volume 303 of *Lecture Notes in Mathematics*, pages 99–110. Springer Berlin Heidelberg, 1972.

[18] E. Köhler, R. H. Möhring, and H. Schilling. Fast point-to-point shortest path computations with arc-flags. In *9th DIMACS Implementation Challenge*, 2006.

[19] I. Nitto and R. Venturini. On compact representations of all-pairs-shortest-path-distance matrices. In P. Ferragina and G. M. Landau, editors, *Combinatorial Pattern Matching, 19th Annual Symposium, CPM 2008, Pisa, Italy, June 18-20, 2008, Proceedings*, volume 5029 of *Lecture Notes in Computer Science*, pages 166–177. Springer, 2008.

[20] M. Pătraşcu. Succincter. In *49th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2008, October 25-28, 2008, Philadelphia, PA, USA*, pages 305–313. IEEE Computer Society, 2008.

[21] M. Pătraşcu and L. Roditty. Distance oracles beyond the Thorup-Zwick bound. *SIAM J. Comput.*, 43(1):300–311, 2014.

[22] L. Roditty. Distance oracles for sparse graphs. In M. Kao, editor, *Encyclopedia of Algorithms*. Springer, 2015.

[23] C. Sommer, E. Verbin, and W. Yu. Distance oracles for sparse graphs. In *50th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 703–712. IEEE, 2009.

[24] M. Thorup and U. Zwick. Compact routing schemes. In *Proceedings of the Thirteenth Annual ACM Symposium on Parallel Algorithms and Architectures*, SPAA '01, pages 1–10, New York, NY, USA, 2001. ACM.

[25] O. Weimann and D. Peleg. A note on exact distance labeling. *Inf. Process. Lett.*, 111(14):671–673, 2011.

[26] C. Wulff-Nilsen. Approximate distance oracles with improved query time. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 539–549.